

**Subject:** Re: FTB workflows  
**From:** Aurélien Bouteiller <bouteill@eecs.utk.edu>  
**Date:** Tue, 17 Feb 2009 16:30:20 -0500  
**To:** cifts@googlegroups.com

Here are the workflows we envisioned regarding FTLA and our HEAT application at UTK.

Workflow #1:

FTLA Application cooperates with job scheduler to recover from node failure

1. Job scheduler listen for CAN\_SURVIVE\_FAILURE events
2. Application job#k notify it can tolerate failures by sending the aforementioned event
3. Job scheduler/resource manager mark job#k as a "resilient" application
4. RAS or Resource manager of application detects failure of node#N and sends a NODE\_FAILURE(N)
5. Job scheduler receive NODE\_FAILURE(N), sort out that node#N is part of job#k
6. Job#k detects node#N is failed. Application computes available recovery actions and computes an estimate of various strategies. It sends out a message CAN\_RECOVER(#new\_nodes: walltime, no\_new\_nodes: walltime).
7. Job scheduler and resource manager exchange messages to determine if they can (or want to) give new nodes to job#k and send out a reply SPARE(node#N+1) or NOSPARE or CHECKPOINT or CLEANUP
8. job#k reacts according to answer, if CHECKPOINT, try to take checkpoint and send CHECKPOINT\_ACK when done.

Possible extensions: notifications of the precision loss to let an autonomic component take a decision about restarting from scratch if precision is too bad after recovery.

Workflow #2:

HEAT application cooperates with filesystem to manage checkpoint data failure

1. HEAT application register for FILE\_LOST events
2. HEAT application stores checkpoints in the PVFS
3. Failures hit PVFS, PVFS sends a notification FILE\_LOST
4. HEAT application takes corrective actions to rebuild missing parts of checkpoint data

Possible extensions: managing checkpoint generated contention by letting file system hint when it is idle and therefore available to checkpoint traffic.

Thanks,  
Aurelien

--

\* Dr. Aurélien Bouteiller  
\* Sr. Research Associate at Innovative Computing Laboratory  
\* University of Tennessee  
\* 1122 Volunteer Boulevard, suite 350  
\* Knoxville, TN 37996  
\* 865 974 6321

-----  
You received this message because you are subscribed to the Google Groups "CIFTS" group.  
To post to this group, send email to [cifts@googlegroups.com](mailto:cifts@googlegroups.com)  
To unsubscribe from this group, send email to [cifts+unsubscribe@googlegroups.com](mailto:cifts+unsubscribe@googlegroups.com)  
For more options, visit this group at <http://groups.google.com/group/cifts?hl=en>  
-----